

Adding Analytical Behavioral Intelligence to Block Storage Layer

andy mills, co-founder/ceo

XLDB 2016
May 24-26, 2016

www.enmotus.com

Topics

- Large scale software defined storage (SDS) challenges
- Need for a fast, efficient storage device virtualization (SDV) layer
- Behavioral analysis and automation of storage devices

Evolving Applications, Stack and Devices

Applications / File System / Objects

Storage Services Layer



Hyperconverged Storage

NUTANIX ATLANTIS

Scaleout Storage Systems

hadoop Spark ceph

Virtual Data Center

openstack KVM vmware ESXi Microsoft Hyper-V docker Xen

Operating System/ Cloud Instance

Microsoft Linux aws Microsoft Azure

....

THE VOID / THE NEED

- Device Agility
- Device Automation
- Device Management
- Remove Long-Tail Latency
- Device Analytics
- New Storage Technology Enablement
- File/Object/OS Independence
- Node Right-Sizing



Flash/NVM/SCM

3DXP MLC SLC

intel SAMSUNG TOSHIBA Micron

HDD

SMR PMR HAMR

SEAGATE WD Western Digital TOSHIBA

Archive (Optical & Tape)

SONY Panasonic

Interfaces and Device Management

NVMe Fabric Connect

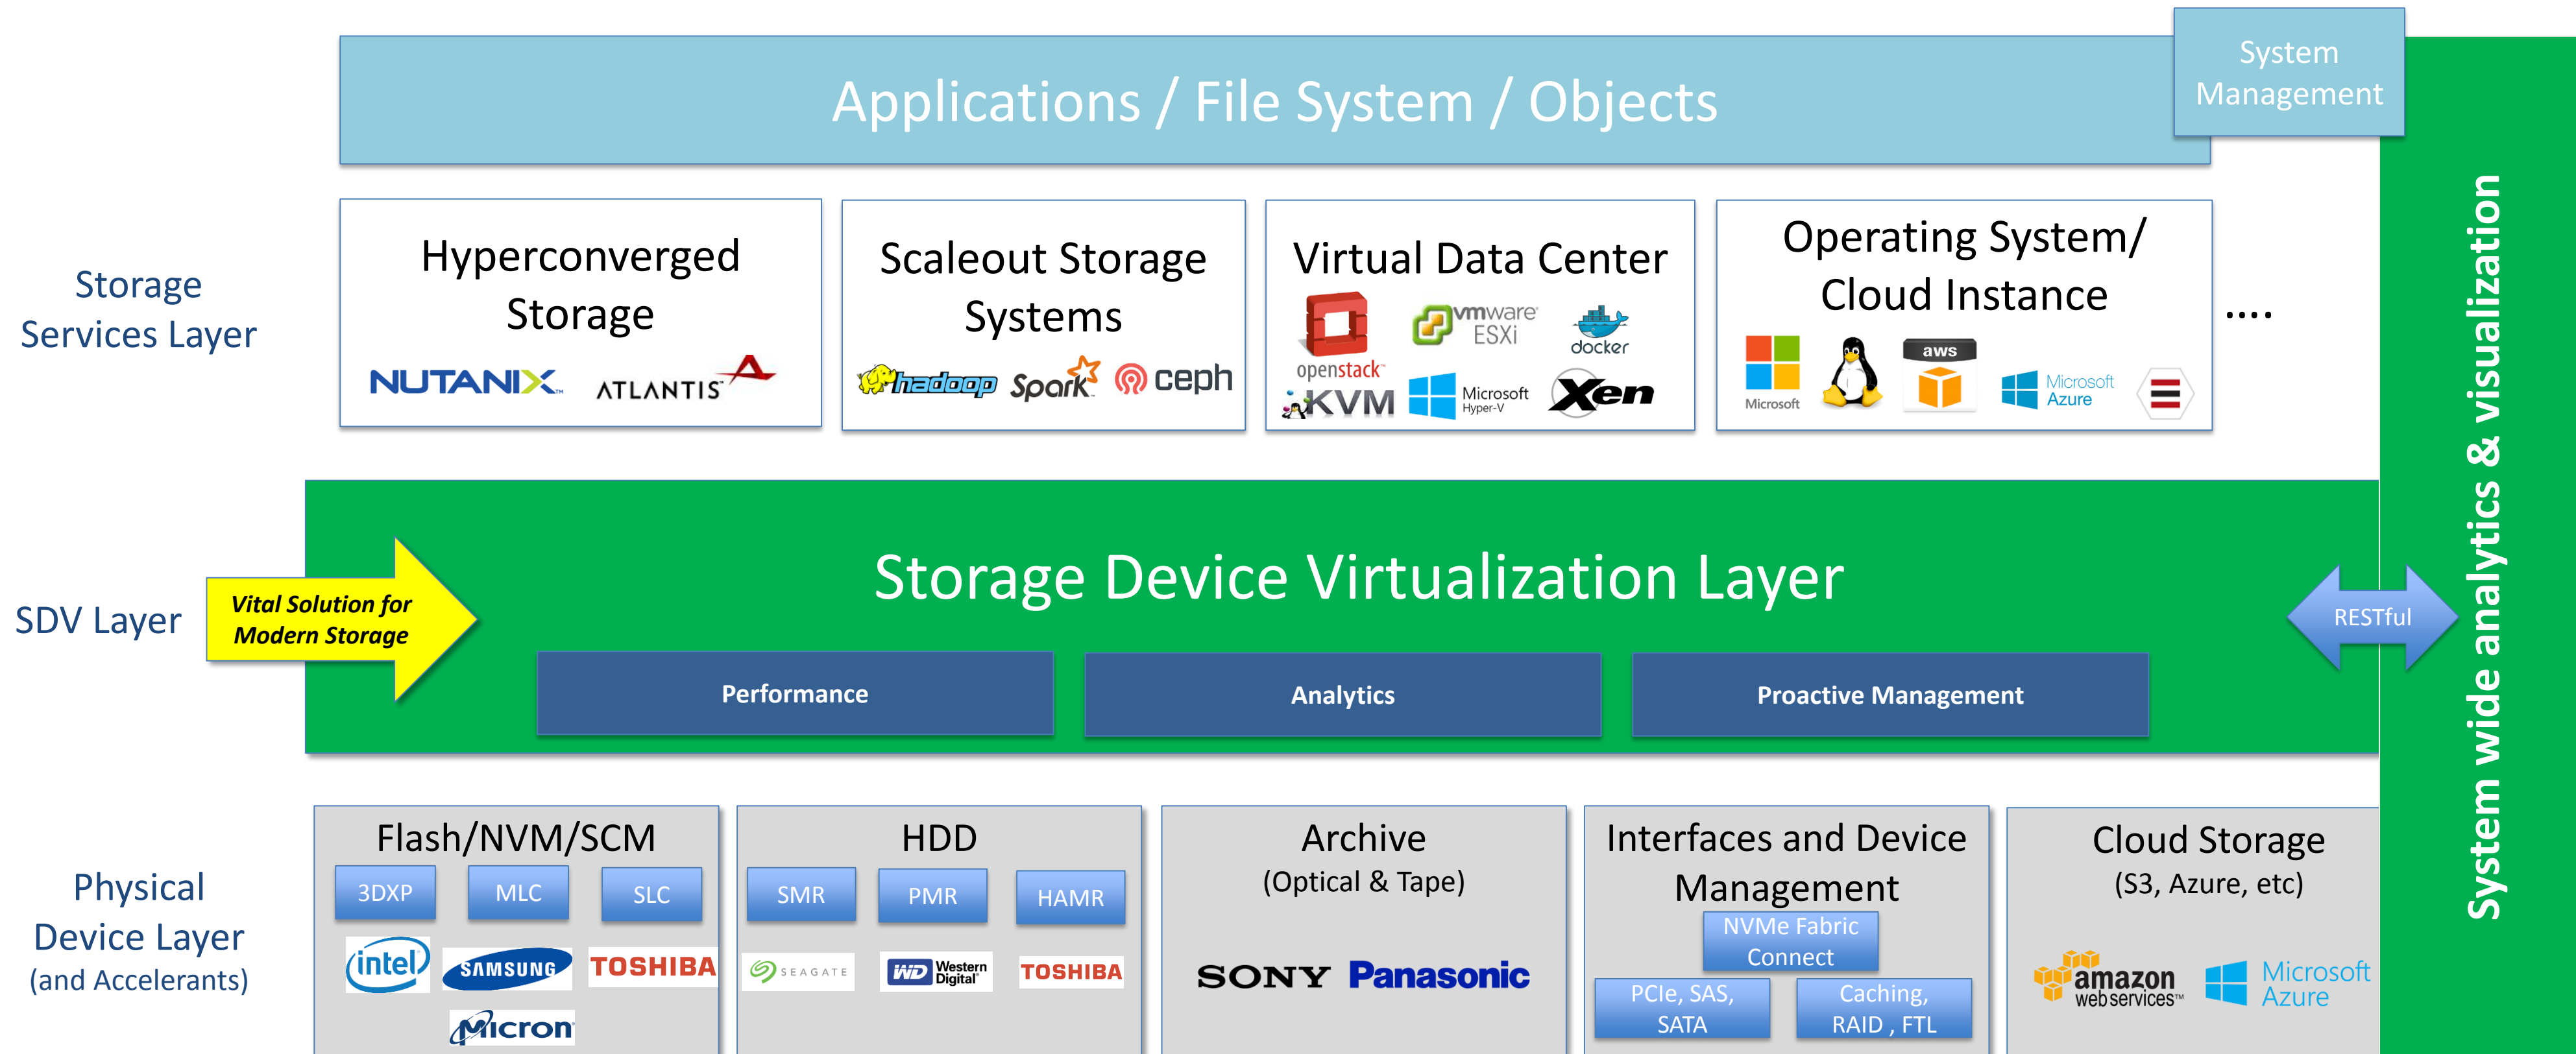
PCIe, SAS, SATA Caching, RAID, FTL

Cloud Storage (S3, Azure, etc)

amazon web services Microsoft Azure

Physical Device Layer (and Accelerants)

Evolving Applications, Stack and Devices




Storage Device Virtualization

- Intelligent software device software layer
 - Behavioral approach to mapping devices to application workloads
 - Autonomous and centralized device management
 - Fast translation i.e. minimal impact IO performance and latency
- Benefits
 - Node level – automatically load balance across RAM, SSD, HDD
 - System level – detect and isolate issues such as long tail latency
 - Central collector – analyze and correct device behavior
- Open - provide APIs via JSON/RESTful protocols
 - Connectors to other tools e.g. Splunk or internal management

Google's Disk for Data Centers

Key Problems Identified	Storage Device Virtualization (SDV)
Balanced application of DRAM/SSD/HDD	Automated, intelligent real time block or memory migration between devices
Move cache from disks to hosts	Choose most appropriate cache media RAM, NVRAM, SSD
Hybrid use of CMR and SMR drives	Automatically map to CMR or SMR (all types) based on detected traffic patterns
Host managed retries to contain tail latency	Manage long tail latency through both active and passive behavioral analysis versus than just simple SMART logs reporting
Capture more performance info to manage tail latency	Combined spatial and temporal statistics can better determine where the origins of tail latency lie and enable better automation of fixes
Flash device behavior with respect to uncorrectable events is problematic	Machine level behavioral analysis can automatically correct problematic devices

SDV-Analytics In Action



Live fast media remapping vs. performance sequence over 20 minutes (Click to animate)

- Visualize in real time how media adapts to application workloads
- Real time or over extended periods of time
- Deep insight trend analysis

Thanks!

